DRAFT 4 FOR DISCUSSION ONLY 6/5/06

Washington State Health Information Infrastructure Advisory Board

Review and Comparison of Architectural Models for Health Information Infrastructure

Executive Summary

[TO BE ADDED LATER]

Background

A key issue for the HIIAB is to develop a recommended architectural model for health information infrastructure. While such architectural choices might seem highly detailed and technical, and therefore beyond the scope of HIIAB discussions, in this case the specific model used will have substantial and overarching implications for the operation, financing, privacy, and overall usefulness of the system. Therefore, the HIIAB is carefully reviewing the characteristics of various alternatives to assist in making the best possible choice. To aid in this process, this document presents a review of the models under consideration and compares their characteristics.

Architectural Models (see animated illustrations in accompanying slide show)

A. Distributed

In the distributed model, all medical record information remains in its current location. A central "record locater service" (RLS) is established in the community containing a list of where records may be found for each person. When the records are needed, the RLS sends record requests to each system on that patient's list, receives the records from each system queried, aggregates the records, and makes them available. After the care is completed, a "pointer" is sent to the RLS indicating that records are now available in that location for retrieval next time the patient's medical records are needed.

B. Central

In the central repository model, the complete lifetime health record (LHR) for each person in the community is stored at a single location, known as a health record bank (HRB). When a patient receives care, his/her LHR is retrieved from the central repository. After the care is completed, a copy of the record of that care is sent electronically to the central repository. If it is a paper record, it can be faxed to the repository and stored as an image (although the information contained in the record would not be available for computer processing).

The system must include state-of-the-art computer security to minimize the possibility of data loss. This security may include a separate server used for searching, which has no phone lines or network connections. In this way, the server used to access individual patient records may be configured without any searching or other data retrieval capabilities beyond what is needed to access a single patient's record.

C. Competitive banking

In the competitive banking model, multiple organizations in the community offer "health record banks" (HRBs) where patients may choose to store their complete lifetime health record (LHR). A central "routing and lookup" service is established that keeps track of which HRB holds the record for each person. When the record is needed for care, a request for the record is sent to the "routing and lookup" server which retrieves the record from the HRB and makes it available. When the care episode is completed, the information generated is sent to the "routing and lookup" server which forwards it to the patient's HRB for storage.

The competitive banking model provides a defined mechanism for linking together HRBs that allows each one to operate independently while being connected to the overall system. To enable this, each HRB must provide interfaces for three functions: 1) retrieval of the patient LHR; 2) input of new information for the patient LHR; and 3) authorized searching of the available LHR records. Within a community functioning as an HRB, either a centralized or decentralized architecture could be used. In the decentralized case, there would need to be a "foreign information repository" to store information sent to the HRB from the outside. Any information about a care episode from outside the HRB area would be stored in this "foreign information repository" and a pointer would be added to the "record locater service" system so that it would be found next time the record was needed.

Note that the centralized community repository (described above) is a special case of the competitive banking model where there is a single bank for the entire community.

Criteria for comparison

The models were compared using a subset of the requirements established by the HIIAB. Certain requirements were excluded because they are unaffected by the architectural model (e.g. the characteristics of the organization that operates the HII system). The HIIAB requirements shown in Table 1 are marked as to relevance to the architectural model.

In addition, the HIIAB has previously identified the importance of flexibility and scalability, as well as the ability of the architecture to incorporate existing systems to the greatest extent possible.

Evaluation and Comparison of the Options

Table 2 shows to what extent each of the three models addresses the operational requirements established by the HIIAB.

1. Lifetime health record (LHR) is available

This is supported by all three models. However, the response time for access to the LHR is likely to be problematic for the distributed model. In order to assemble the LHR, the distributed model must issue a substantial number of secondary queries to those systems that have the actual information. The overall response time will be equivalent to the slowest response time of any of these secondary queries, since the complete record cannot be assembled until all the information has been returned. The distributed model also is susceptible to incomplete information if any of the systems with patient information are not operational.

2. Patients control access

All three models can support this requirement.

3. Error correction

This is supported except in the case of the distributed model. Since the information is held by others, action taken in response to correction requests depends on the policies and procedures of each organization holding patient information.

4. LHR information transfer

All three models can support this requirement.

5. Searching data for public health and medical research

This is supported except in the case of the distributed model. Since the information is distributed, it must be assembled before it can be searched. This requires sequential assembly of every record to be searched, a time-consuming and resource-intensive process that will make searching infeasible.

In the other models, searching of the available repository is straightforward. A mechanism for aggregating search results of multiple communities or HRBs would need to be established to accommodate queries that exceeded the scope of a given repository. Such a mechanism would be easy to implement (e.g. the search could be done separately in each locality with the results transmitted to a central point for aggregation).

6. Information is associated with the correct person

This is supported in the central repository and competitive banking models. When new information is received, it is added to the appropriate person's record. If there are any doubts about which person the information applies to, further investigation can be done (including contacting the patient).

In the distributed model, the identification of the information must occur "on the fly" as the information is assembled from the various sources. This constraint requires automatic, rapid matching that is normally biased toward avoiding combining information incorrectly. There is not sufficient time during retrievals to allow for human intervention. Therefore, this is a challenging problem for the distributed model.

7. Information is transmitted using standards

All three models require this, but it is a more difficult requirement for the distributed model because of the large number of transactions that must be handled in real-time when a record is being assembled. Any problems with standard transactions in the distributed model can result in an incomplete patient record. This issue can be overcome with careful testing and monitoring of transactions from every possible source.

8. Healthcare providers can use any system they choose if it can transmit data using standards

This condition is met by all three models. However, the distributed model requires additional capabilities for each healthcare provider system -- namely, that it be able to respond rapidly at any time to queries for information. In the distributed model,

healthcare providers with systems that do not support these query capabilities would not be able to effectively share their information. The other models only require that new information be sent to a health record bank when it is created; no capability for responses to queries is needed.

10. Persons may access a record of all accesses to their own lifetime health record

All three models support this requirement.

11. Complies with privacy and security laws and regulations

All three models can support this requirement, but it is more difficult in the distributed model because of the involvement of multiple data sources in each query. Complex agreements among all the parties may be needed to enable "on the fly" information sharing. These agreements can be simplified if the request for information is clearly done with patient consent, but still require extensive negotiation with all institutions and organizations involved. In the central and health record banking models, the information is always provided at the explicit request of the patient, and the source organization no longer controls it after is sent for inclusion in the patient's lifetime health record.

14. Financial sustainability

The centralized and health record banking models are potentially financially sustainable through fees to patients (and possibly other stakeholders). Sustainability is aided by the lower costs compared to the distributed model (due to the smaller number of transactions and reduced complexity). In the distributed model, a query is generated to the source holder of patient information each and every time the person's lifetime health record is needed. Compare this to a single transmission of that same data to a centralized or health record banking facility. Between approximately 20 and 100 secondary queries are likely to be needed each time a record is requested for the average patient with the distributed model,[1] versus zero for the other models. Enabling this large volume of communications traffic will require substantial costs for additional hardware, software, communications bandwidth, and technical support personnel. All these additional costs will make achieving financial sustainability much more difficult.

15. The system is highly reliable and highly available

This will be much more challenging with the distributed model because of the substantial communications traffic and the dependence on thousands of secondary

---

[1] Assume the average person has three medical encounters/year for a 70-year lifetime (210 total -- round to 200 for simplicity). An optimistic estimate of the average number of sources needed for a typical query would be half these encounters (100) (optimistic because the distribution of queries is non-linear, since the number of health care encounters increases with age, so most queries would occur later in life when there are more sources of information). Then the number of secondary queries would vary depending on what percentage of these encounters were at places that had previously been visited. A very high average for repeat encounters would be about 80%, which would result 20 secondary queries. A low average would be 0%, resulting in 100 secondary queries. Therefore, the range of 20-100 secondary queries is used. Note that this also assumes that each encounter generates information at only one source -- in reality, it is likely to be higher which would further increase the average number of secondary queries.

systems being available for queries. The centralized and health record banking models can be highly reliable and highly available by providing appropriate backup systems.

17. Accommodates existing infrastructure

The competitive banking model also allows maximal use of existing systems. Organizations that have already implemented EHR systems may elect to use them to provide health record banking, both for their own patients and others in the community who may wish to have their records stored. This should ease the cost burden of establishing health record banks, since the existing EHRs represent a sunk cost, and the marginal cost of adding additional records to such systems is small.

18. Scalability

The distributed model is challenging even at a small scale, and becomes increasingly difficult as it grows to encompass more information sources. It has not yet been shown to be workable for populations beyond 500,000. The centralized model probably works best for populations in the 1-10 million range, beyond which the economies of scale are probably outweighed by the large data storage requirements and the inherent risks of having health records for such a large population be dependent on a single facility (even with backup). The competitive health record banking model can be scaled arbitrarily and will still work well using, for example, an indexing system similar to the DNS (domain name server) concept that routes traffic on the Internet (as long as the component banks do not exceed the size limits for their own internal architectural models). In addition, this latter model is indifferent to the details of operations within the component banks, therefore allowing a multiplicity of solutions for various communities and regions. This scalability makes the competitive banking model an attractive policy option.

Conclusion

The competitive banking model provides the greatest degree of both functionality and flexibility, particularly in the context of a statewide solution. Because several communities have already embarked on implementations of HII, it is highly desirable to provide the ability to link those existing systems into a coherent whole. The competitive banking model even allows communities to choose either a central, distributed, or competitive banking model internally, provided they implement the two required interfaces to link statewide (i.e. access to patient records and receipt of new information for their patients from other locations). It therefore provides maximum choice for communities while simultaneously defining a scalable infrastructure that can operate statewide and connect to the rest of the nation (and beyond).

Table 1. Relevance of Proposed HIIAB Requirements to Architectural Model

Functions

**1. The substantive lifetime health record (LHR*) of each person from all sources (with each source identified) is available to authorized users when/where needed, but unavailable otherwise

2. Participation in the HII system is voluntary and available to all

**3. Each person controls access to each portion of their own LHR

**4. Incomplete information or errors in LHR information can be addressed by authorized users

**5. All or part of a person's LHR information may be transferred securely and electronically to authorized recipients at their request

**6. With voluntary patient authorization, LHR information may be made available for public health (including biosurveillance) and medical learning

**7. All information maintained by the system is reliably associated with the correct person

**8. LHR information is transmitted electronically using national standards whenever available (and system standards when not)

**9. Healthcare providers and organizations are able to use whatever information system(s) they choose, provided they can transmit and receive information using designated standards

Privacy, Confidentiality, and Security

10. All users are reliably authenticated

**11. Persons may access a record of all accesses to their own LHR information

**12. The HII system complies with all relevant privacy and security statutes and regulations

13. HII system security is up-to-date and reviewed periodically.

Organizational & Financial

14. A trusted organization operates HII system

**15. The HII system is financially sustainable

**16. The HII system is highly available and highly reliable

**17. The HII accommodates the use of existing infrastructure

**18. The HII system is scalable to accommodate all residents of Washington State

Table II. Ability of Models to Address Requirements

Requirement	Distributed	Central	Competitive Banking
LHR available (1)	YES, but response time may be a problem & requires national interoperability	YES	YES
Patients Control Access (3)	YES	YES	YES
Error Correction (4)	Depends on policies of information holders	YES	YES
LHR info transfer (5)	YES	YES	YES
Public health/research (6)	NO	YES	YES
Info associated with correct person (7)	Challenging	YES	YES
Use of Standards (8)	YES	YES	YES
Use any EHR (9)	YES (but extra query capabilities needed)	YES	YES
Audit trails available to patients (11)	YES	YES	YES
Complies with privacy and security (12)	YES, with extensive policy development	YES	YES
Financially sustainable (15)	Difficult due to high cost & complexity	Potentially	Potentially
Highly available & reliable (16)	Difficult due to high complexity	YES	YES
Accommodates existing infrastructure (17)	YES; needs extensive new communications	YES; needs repository	YES; only needs router
Scalable (18)	Response time & complexity increases with size	YES	YES

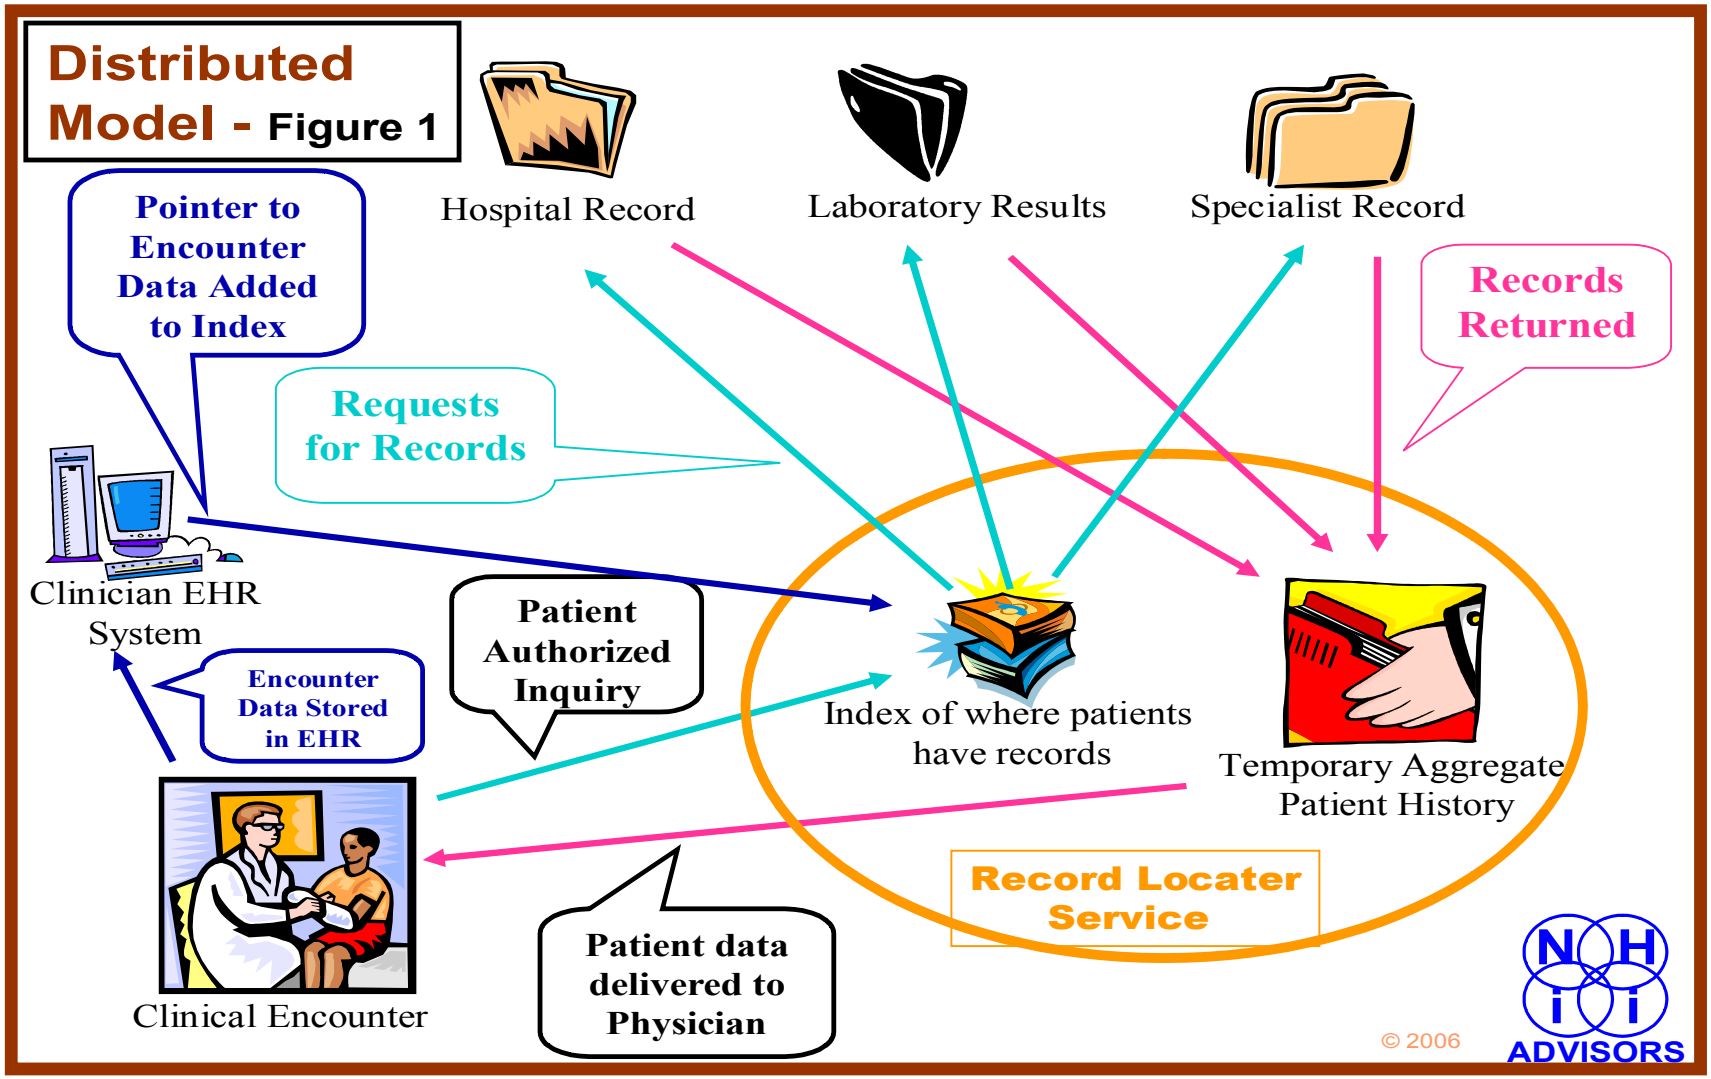
Distributed Model - Figure 1
Hospital Record
Laboratory Results
Specialist Record
Pointer to Encounter Data Added to Index
Requests for Records
Records Returned
Clinician EHR System
Patient Authorized Inquiry
Encounter Data Stored in EHR
Index of where patients have records
Temporary Aggregate Patient History
Record Locater Service
Clinical Encounter
Patient data delivered to Physician
© 2006
NHII ADVISORS

Centralized Model - Figure 2

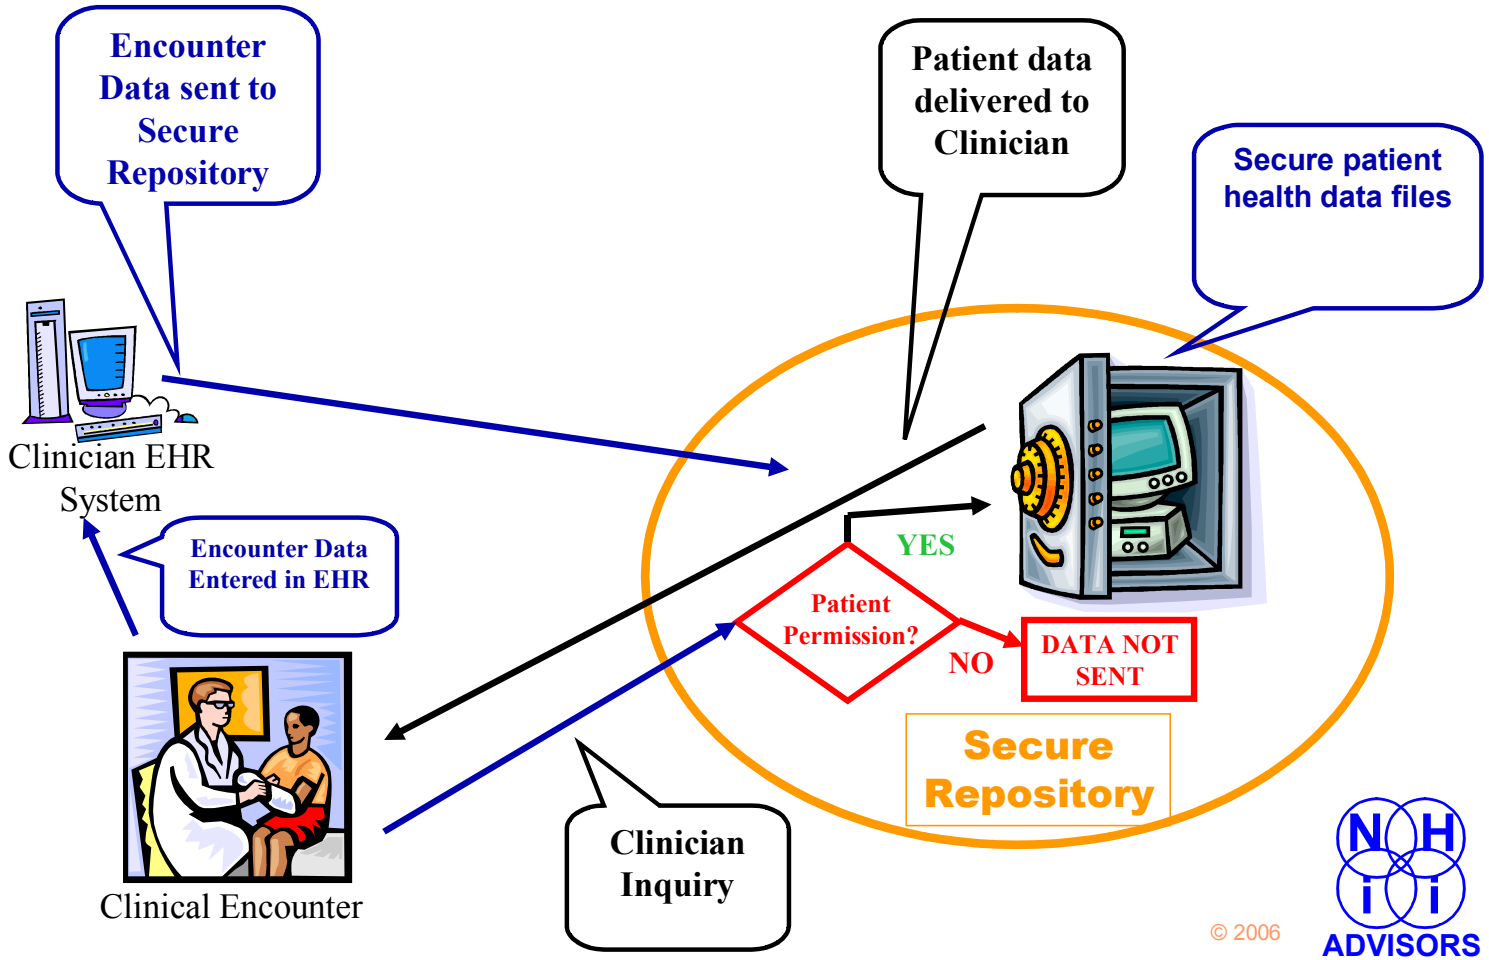

Competitive Banking Model - Figure 3